von Neumann and Morgenstern's theory

"Expected utility"

Objects of choice called *lotteries.* Simplification: Each has only two possible, mutually exclusive outcomes. Notation: $L(x, y, \pi)$ means:

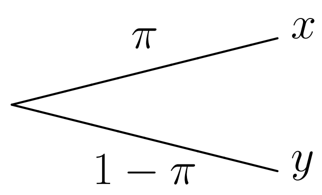

(The $L()$ notation means: The first two arguments are outcomes. Then comes the probability (here: $\pi \in [0, 1]$) of the outcome mentioned first (here: x).)

Axiom C.2 (D&D, p. 45) says that an individual is able to compare and choose between such stochastic variables, and that preferences are transitive. Axiom C.3 says that preferences are continuous. Assumptions like C.2 and C.3 are known from standard consumer theory. Axiom C.1 says that only the probability distribution matters.

Axioms C.4–C.7 specific to preferences over *lotteries.* The theory assumes axioms C.1–C.7 hold for the preferences of one individual. Using the theory, we usually assume it holds for all individuals, but their preferences may vary within the restrictions given by the theory.

Axiom C.4 *Independence:*

Let x, y and z be outcomes of lotteries. In fact, x, y, and/or z could be new lotteries. Assume $y \succsim z$. Then

$$L(x, y, \pi) \succsim L(x, z, \pi).$$

Axiom C.5

Among all lotteries (and outcomes), there exists one best lottery, b, and one worst, w, with $b \succ w$.

Axiom C.6

If $x \succ y \succ z$, then there exists a unique π such that

$$y \sim L(x, z, \pi).$$

(Not obvious. What about life and death?)

Axiom C.7

Assume $x \succ y$. Then

$$L(x, y, \pi_1) \succ L(x, y, \pi_2) \Leftrightarrow \pi_1 > \pi_2,$$

(Actually: None of axioms are obvious.)

Derivation of theorem of expected utility

With reference to b and w, for all lotteries and outcomes z, define a function $\pi()$ such that

$$z \sim L(b, w, \pi(z)).$$

This probability exists for all z by axiom C.6. By axiom C.7 it is unique and can be used to rank outcomes, since $\pi(x) > \pi(y) \Rightarrow x \succ y$. Thus $\pi()$ is a kind of utility function. Will prove it has the expected utility property: The utility of a lottery is the expected utility from its outcomes.

Digression: A *utility function* for a person assigns a real number to any object of choice, such that a higher number is given to a preferred object, and equal numbers are given when the person is indifferent between the objects. If x and y are money outcomes or otherwise quantities of a (scalar) good, and there is no satiation, then π is an increasing function.

The expected utility property

Consider a lottery $L(x, y, \pi)$, which means:

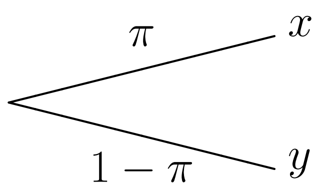

When $x \sim L(b, w, \pi_x)$ and $y \sim L(b, w, \pi_y)$, then there will be indifference between $L(x, y, \pi)$ and each of these two:

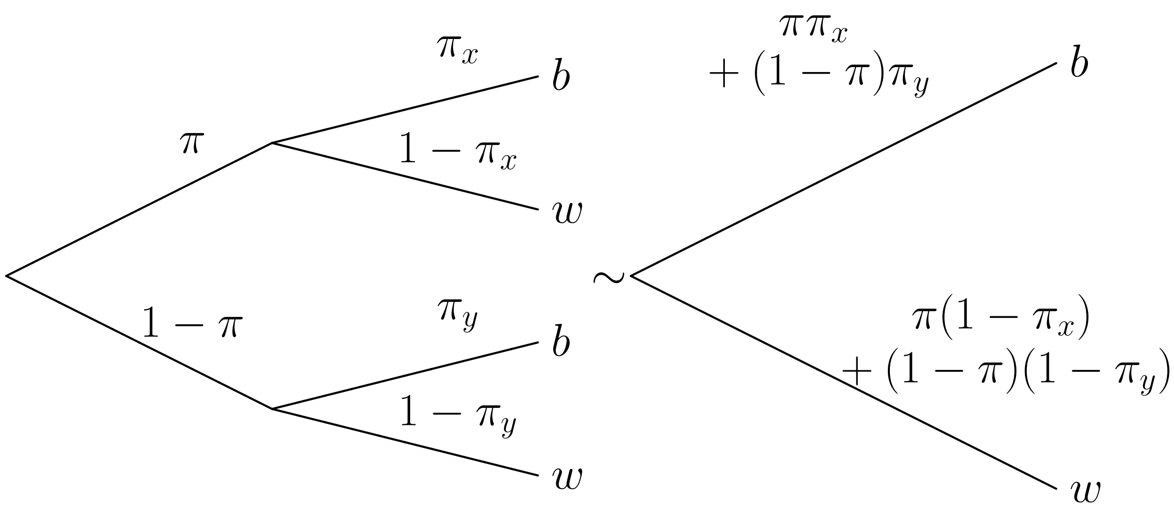

So that

$$L(x, y, \pi) \sim L(b, w, \pi\pi_x + (1 - \pi)\pi_y).$$

Thus the "utility" of $L(x, y, \pi)$ is $\pi\pi_x + (1 - \pi)\pi_y$.

Write $\tilde{X}$ for the stochastic variable with outcomes x_1 with probability π_1 and x_2 with probability $\pi_2 = 1 - \pi_1$. The utility expression $\pi\pi_1 + (1-\pi)\pi_2$ can be interpreted as $E[\pi(\tilde{X})]$, which explains why it is called *expected utility.*

Notation: Usually the letter U is chosen for the utility function instead of π, and expected utility of $\tilde{X}$ is written $E[U(\tilde{X})]$.

Possible to extend to ordering of lotteries of more than two outcomes,

$$E[U(\tilde{X})] = \sum_{s=1}^{S} \pi_s U(x_s),$$

even to a continuous probability distribution,

$$E[U(\tilde{X})] = \int_{-\infty}^{\infty} U(x) f(x) dx.$$

Will not look at this more formally.

Criticism of vN-M expected utility

- Some experiments indicate that many people's behavior in some situation contradicts expected utility maximization.
- Exist alternative theories, in particular generalizations (alternative theories in which expected utility appears as one special case).
- Nevertheless much used in theoretical work on decisions under uncertainty.

Example of when vN-M may not work

- Suppose every consumption level below 5 is very bad.
- Suppose, e.g., that $U(4) = -10, U(6) = 1, U(8) = 4, U(10) = 5$.
- Then $E[U(L(4, 10, 0.1))] = 0.1 \cdot (-10) + 0.9 \cdot 5 = 3.5$, while $E[U(L(6, 8, 0.1))] = 0.1 \cdot 1 + 0.9 \cdot 4 = 4.6$.
- But even with the huge drop in U level when consumption drops below 5, one will prefer the first of these two alternatives (the lottery $L(4, 10, \pi)$) to the other ($L(6, 8, \pi)$) as soon as π drops below 1/12.
- If one outcome is so bad that someone will avoid it *any* cost, even when its probability is very low, then that person's behavior contradicts the vN-M theory.
- In particular, axiom C.6 is contradicted.

Allais paradox

Behavior at odds with vN-M theory, observed by French economist Maurice Allais. Consider the following lotteries:

- $L^3 = L(10000, 0, 1)$
- $L^4 = L(15000, 0, 0.9)$
- $L^1 = L(10000, 0, 0.1) = L(L^3, 0, 0.1)$
- $L^2 = L(15000, 0, 0.09) = L(L^4, 0, 0.1)$

People asked to rank L^1 versus L^2 often choose $L^2 \succ L^1$. (Probability of winning is just slightly less, while prize is 50 percent bigger.) But when these same people are asked to rank L^3 versus L^4, they often choose $L^3 \succ L^4$. (With strong enough risk aversion, the drop in probability from 1 to 0.9 is enough to outweigh the gain in the prize.) Is this consistent with the vN-M axioms?

Uniqueness of U function?

Given a vN-M preference ordering of one individual, have now shown we can find a U function such that

$$\tilde{X} \succ \tilde{Y} \text{ if and only if } E[U(\tilde{X})] > E[U(\tilde{Y})].$$

Considering one individual, we ask: Is U unique? No, depends on b and w, but preferences between $\tilde{X}$ and $\tilde{Y}$ do not.

Define an increasing linear transformation of U,

$$V(x) \equiv c_1 U(x) + c_0,$$

where $c_1 > 0$ and c_0 are constants. This represents the preferences of the same individual equally well since

$$E[V(\tilde{X})] = c_1 E[U(\tilde{X})] + c_0$$

for all X, so that a higher $E[U(\tilde{X})]$ gives a higher $E[V(\tilde{X})]$, and vice versa.

But *not* possible to do similar replacement of U with any *non*-linear transformation of U (as opposed to ordinal utility functions for usual commodities). For instance, $E\{\ln[U(\tilde{X})]\}$ does not necessarily increase when $E[U(\tilde{X})]$ increases. So $\ln[U()]$ cannot be used to represent the same preferences as $U()$.

Risk aversion

For those preference orderings which (i.e., for those individuals who) satisfy the seven axioms, define *risk aversion*.

Compare a lottery $\tilde{Y} = L(a, b, \pi)$ (where a, b are fixed monetary outcomes) with receiving $E(\tilde{Y}) = \pi a + (1 - \pi)b$ for sure. Whether the lottery, $\tilde{Y}$, or its expectation, $E(\tilde{Y})$, is preferred, depends on the curvature of U:

- If U is linear, then $U[E(\tilde{Y})] = E[U(\tilde{Y})]$, and one is indifferent between lottery and its expectation. One is called *risk neutral*.
- If U is concave, then $U[E(\tilde{Y})] \geq E[U(\tilde{Y})]$, and one prefers the expectation. One is called *risk averse*.
- If U is convex, then $U[E(\tilde{Y})] \leq E[U(\tilde{Y})]$, and one prefers the lottery. One is called *risk attracted*.

The inequalities follow from *Jensen's inequality* (see Sydsæter, Strøm and Berck, equations 7.16–7.17 & 33.19, or D&D, p.63). If U is *strictly* concave or convex, the inequalities are strict, except if $\tilde{Y}$ is constant with probability one.

Quite possible that many have U functions which are neither everywhere linear, everywhere concave, nor everywhere convex. Then one does not fall into one of the three categories.

Assume risk aversion

(Risk aversion does *not* follow from the seven axioms.)

- Most common behavior in economic transactions.
- Explains the existence of insurance markets.
- But what about money games? Expected net result always negative, so a risk-averse should not participate. Cannot be explained by theories taught in this course.
- Some of our theories will collapse if someone is risk neutral or risk attracted. Those will take all risk in equlibrium. Does not happen.

How measure risk aversion?

- Natural candidate: $-U''(y)$. (Why minus sign?)
- Varies with the argument, e.g., high y may give lower $-U''(y)$.
- Is $U()$ twice differentiable? Assume yes.
- But: The magnitude $-U''(y)$ is not preserved if $c_1U() + c_0$ replaces $U()$.
- Use instead:
    - $-U''(y)/U'(y)$ measures *absolute risk aversion.*
    - $-U''(y)y/U'(y)$ measures *relative risk aversion.*
- In general, these also vary with the argument, y.

Arrow-Pratt measures of risk aversion

- Will introduce the concept *risk premium*, related to expected utility. This concerns a situation in which we have specified the complete, uncertain consumption (or income or wealth) which is the argument of the (expected) utility function.
- (One could also say that the agents in a CAPM world require a risk premium for undertaking a risky investment, but in that setting, the investment is seen as a small addition to the rest of their wealth, which is diversified. This would be a different concept of a risk premium.)
- Will also say more about the two measures of risk aversion.
- Will show on next page: For small risks, $R_A(y) \equiv -U''(y)/U'(y)$ meaures how much compensation a person demands for taking the risk. Called the Arrow-Pratt measure of absolute risk aversion.
- $R_R(y) \equiv -U''(y) \cdot y/U'(y)$ is called the Arrow-Pratt measure of relative risk aversion.
- Consider the following case (somewhat more general than D&D, sect. 4.3.1):
  - The wealth Y is non-stochastic.
  - A lottery $\tilde{Z}$ has expectation $E(\tilde{Z}) = 0$.
- For a person with utility function $U()$ and inital wealth Y, define the *risk premium* Π associated with the lottery $\tilde{Z}$ by
$$E[U(Y + \tilde{Z})] = U(Y - \Pi).$$
- Will show the relation between Π and absolute risk aversion.

Risk premium is proportional to risk aversion

(The result holds approximately, for small lotteries.)

$$E[U(Y + \tilde{Z})] = U(Y - \Pi).$$

Take quadratic approximations (second-order Taylor series). (Sorry, the math course this spring will not cover Taylor series until 25 April.)

LHS:

$$U(Y + z) \approx U(Y) + zU'(Y) + \frac{1}{2}z^2U''(Y)$$

which implies

$$E[U(Y + \tilde{Z}]) \approx U(Y) + \frac{1}{2}E(\tilde{Z}^2)U''(Y).$$

RHS:

$$U(Y - \Pi) \approx U(Y) - \Pi U'(Y) + \frac{1}{2}\Pi^2U''(Y).$$

Use the notation $\sigma_z^2 \equiv \text{var}(\tilde{Z})$. This is $= E(\tilde{Z}^2)$ since $E(\tilde{Z}) = 0$.

Since Π is small, Π^2 is very small. Thus the last term on the RHS is very small, and we will neglect it. Then we are left with:

$$\frac{1}{2}\sigma_z^2U''(Y) \approx -\Pi U'(Y)$$

which implies the promised result:

$$\Pi \approx -\frac{U''(Y)}{U'(Y)} \cdot \frac{1}{2}\sigma_z^2.$$

The U function: Forms which are often used

- Some theoretical results can be derived without specifying form of U.
- Other results hold for specific classes of U functions.
- *Constant absolute risk aversion (CARA)* holds for $U(y) \equiv -e^{-ay}$, with $R_A(y) = a$.
- *Constant relative risk aversion (CRRA)* holds for $U(y) \equiv \frac{1}{1-g}y^{1-g}$, with $R_R(y) = g$.
- (Exercise: Verify these two claims. (a, g are constants.) Determine what are the permissible ranges for y, a and g, given that functions should be well defined, increasing, and concave.)
- Essentially, these are the only functions with CARA and CRRA, respectively, apart from CRRA with $R_R(y) = 1$.
- (Any constant can be added to the functions, and any positive constant can be multiplied with them.)
- $R_R(y) \equiv 1$ is obtained with $U(y) \equiv \ln(y)$.
- Another much used form: $U(y) = -ay^2 + by + c$, quadratic utility. Easy for calculations, U' linear.
- (What are permissible ranges, given that U should be concave? Hint: There is a minus sign in front of a.)
- Quadratic U has increasing $R_A(y)$ (Verify!), perhaps less reasonable.
- (What happens for this U function when $y > b/2a$? Is this reasonable?)